

Privacy-Preserving Group Data Access via Stateless Oblivious RAM Simulation

Michael T. Goodrich
Dept. of Computer Science
University of California, Irvine
<http://www.ics.uci.edu/~goodrich/>

Michael Mitzenmacher
Dept. of Computer Science
Harvard University
<http://www.eecs.harvard.edu/~michaelm/>

Olga Ohrimenko
Dept. of Computer Science
Brown University
<http://www.cs.brown.edu/~olya/>

Roberto Tamassia
Dept. of Computer Science
Brown University
<http://www.cs.brown.edu/~rt/>

Abstract

We study the problem of providing privacy-preserving access to an outsourced honest-but-curious data repository for a group of trusted users. We show that such privacy-preserving data access is possible using a combination of probabilistic encryption, which directly hides data values, and stateless oblivious RAM simulation, which hides the pattern of data accesses. We give simulations that have only an $O(\log n)$ amortized time overhead for simulating a RAM algorithm, \mathcal{A} , that has a memory of size n , using a scheme that is data-oblivious with very high probability assuming the simulation has access to a private workspace of size $O(n^\nu)$, for any given fixed constant $\nu > 0$. This simulation makes use of pseudorandom hash functions and is based on a novel hierarchy of cuckoo hash tables that all share a common stash. We also provide results from an experimental simulation of this scheme, showing its practicality. In addition, in a result that may be of some theoretical interest, we also show that one can eliminate the dependence on pseudorandom hash functions in our simulation while having the overhead rise to be $O(\log^2 n)$.

1 Introduction

Companies offering outsourced data storage services are defining a growing industry, with competitors that include Amazon, Google, and Microsoft, which are providing outsourced data repositories for individual or corporate users, with prices that amount to pennies per gigabyte stored.

Clearly, the customers of such cloud computing services have an interest in security and privacy, particularly for proprietary data. As a recognition of this interest, we note that, as of November 2010, the Amazon S3 and Microsoft Azure cloud platform have achieved ISO 27001 certification and Google's cloud computing service has SAS70 certification. In spite of these certifications, the companies that provide outsourced data services nevertheless often have commercial interests in learning information about their customers' data. Thus, the users of such systems should also consider technological solutions for maintaining the privacy of their outsourced data in addition to the assurances that come from certifications and formal audits.

Of course, a key component for users to maintain the privacy of their data is for them to store their data in encrypted form, e.g., using a group key known only to the set of users. Simply encrypting the group's data is not sufficient to achieve privacy, however, since information about the data may be leaked by the pattern in which the users access it. For example, at the Oakland 2010 conference, Chen *et al.* [8] show that highly sensitive data, such as financial and health information, can be inferred from access patterns at popular financial and health web sites even if the contents of those communications are encrypted.

1.1 Group Access to Outsourced Data

In this paper, we are interested in technological solutions to the problem of protecting the privacy of a group's data accesses to an outsourced data storage facility. In this framework, we assume that a trusted

group, G , of users shares a group key, K , with which they encrypt all their shared data that is stored at a semi-trusted data outsourcer, Bob. Furthermore, we assume that the users access their data according to a public indexing scheme, which Bob knows; hence, we can model Bob’s memory, M , as in the standard RAM model (e.g., see [1, 9, 14, 18]).

Each time a user, Alice, in G , accesses Bob’s memory, she specifies an index i , and Bob responds with $C = M[i]$. Alice then performs the following (atomic) sequence of operations:

1. She decrypts C using K , producing the plaintext value, $P = D_K(C)$, that was stored in encrypted form at index i by Bob.
2. She optionally changes the value of P , depending on the computation she is performing, producing the plaintext value, P' .
3. She encrypts P' using a probabilistic encryption scheme based on K , producing ciphertext $C' = E_K(P')$.
4. She returns C' to Bob for him to store back in his memory at index i ; that is, she directs Bob to assign $M[i] \leftarrow C'$.

By using a probabilistic encryption scheme, the users in the group G ensure that Bob is computationally unable to determine the plaintext of any memory cell from that cell’s contents alone. Also, it is unfeasible for Bob to determine whether two memory cells store encryptions of the same plaintext.

1.2 Stateless Oblivious RAM Simulation

In addition to using probabilistic encryption, the users in the group G also need to hide their data access patterns from Bob, so as to avoid inadvertent information leaks. To facilitate such information hiding, we formulate the privacy objective of the users in G in terms of the *stateless oblivious RAM simulation* problem.

In this framework, we model the group G as a single user, Alice, who has a register holding the key K and a CPU with a private cache. Alice’s interactions with Bob occur in discrete *episodes* in which she reads and writes a set of cells in his memory, using probabilistic encryption, as described above, to hide data contents. Alice’s cache may be used as a private workspace during any episode, but it cannot store any information from one episode to the next. This requirement is meant to model the fact that Alice is

representing a group of users who do not communicate outside of their shared access to Bob’s memory. That is, each episode could model a consecutive set of accesses from different users in the group G . Moreover, this requirement is what makes this framework “stateless,” in that no state can be carried from one episode to the next (other than the state that is maintained by Bob).

To allow the group of users, which we model by the stateless Alice, to perform arbitrary computations on the data they share and outsource to Bob, we assume that Alice is simulating a RAM computation. We also assume the service provider, Bob, is trying to learn as much as possible about the contents of Alice’s data from the sequence and location of all of Alice’s memory accesses. As mentioned above, however, he cannot see the content of what is read or written (since it is probabilistically encrypted). Moreover, Bob has no access to Alice’s private cache. Bob is assumed to be an *honest-but-curious* adversary [12], in that he correctly performs all protocols and does not tamper with data.

We say that Alice’s sequence of memory accesses is *data-oblivious* if the distribution of this sequence depends only on n , the size of the memory used by the RAM algorithm she is simulating, m , the size of her private cache, and the length of the access sequence itself. In particular, the distribution of Alice’s memory accesses should be independent of the data values in the input. Put another way, this definition means that $\Pr(S | M)$, the probability that Bob sees an access sequence, S , conditioned on a specific configuration of his memory, M , satisfies

$$\Pr(S | M) = \Pr(S | M'),$$

for any memory configuration $M' \neq M$ such that $|M'| = |M|$.

Examples of data-oblivious access sequences for an array, A , of size n , in Bob’s memory, include the following:

- Scanning A from beginning to end, accessing each item exactly once, for instance, to compute the minimum value in A , which is then stored in $A[1]$.
- Simulating a Boolean circuit, \mathcal{C} , with its inputs taken in order from the bits of A .
- Accessing the cells of A according to a random hash function, $h(i)$, as $A[h(1)]$, $A[h(2)]$, \dots , $A[h(n)]$, or random permutation, $\pi(i)$, as $A[\pi(1)]$, $A[\pi(2)]$, \dots , $A[\pi(n)]$.

Examples of computations on A that would *not* be data-oblivious include the following:

- Scanning A from beginning to end, accessing each item exactly once, to compute the index i of the minimum value in A , and then reading $A[i]$ and writing it to $A[1]$.
- Using a standard heap-sort, merge-sort, or quick-sort algorithm to sort A . (None of these well-known algorithms is data-oblivious.)
- Using values in A as indices for a hash table, T , and accessing them as $T[h(A[1])]$, $T[h(A[2])]$, \dots , $T[h(A[n])]$, where h is a random hash function. For example, consider what happens if the values in A are all equal and how unlikely the resulting collision in T would be.

Note that this last example access pattern actually would be data-oblivious if the elements in A were always guaranteed to be distinct, assuming the random hash function, h , satisfies the standard assumptions of the random oracle model (e.g., see [6]).

1.3 Related Prior Results

Data-oblivious sorting is a fundamental problem (e.g., see Knuth [19]), with deterministic schemes giving rise to sorting networks, such as the impractical $O(n \log n)$ AKS network [2, 3, 25, 29] as well as practical, but theoretically-suboptimal, sorting networks [20, 28]. Randomized data-oblivious sorting algorithms running in $O(n \log n)$ time and succeeding with high probability¹ are studied by Leighton and Plaxton [21] and Goodrich [15]. In addition, data-oblivious sorting is finding applications to privacy-preserving secure multi-party computations [30], and it is used in all the known oblivious RAM simulation schemes (including the ones in this paper).

In early work on the topic of oblivious simulation, Pippenger and Fischer [27] show that one can simulate a Turing machine computation of length n with an oblivious Turing machine computation of length $O(n \log n)$, that is, they achieve an amortized $O(\log n)$ time and space overhead for this oblivious simulation.

More recently, Goldreich and Ostrovsky [13] show that a RAM computation using space n can be simulated with an oblivious RAM with an amortized time overhead of $O(\log^3 n)$ per step of the original RAM algorithm and space overhead of $O(\log n)$. Goodrich and Mitzenmacher [16] improve this result by showing that any RAM algorithm, \mathcal{A} , can be simulated in a data-oblivious fashion, with very high probability, in

¹In this paper, we take the phrase “with very high probability” to mean that the probability is at least $1 - O(1/n^d)$, for any given fixed constant $d \geq 1$.

an outsourced memory so that each memory access performed by \mathcal{A} has a time overhead of $O(\log^2 n)$, assuming Alice’s private cache has size $O(1)$. Their scheme has a space overhead of $O(1)$. Incidentally, in the recent CRYPTO 2010 conference, Pinkas and Reinman [26] also claim an oblivious RAM simulation result having a time overhead of $O(\log^2 n)$, but there is a flaw in this version of their scheme².

In addition to these stateless oblivious RAM simulation schemes, Williams and Sion [31] show how to simulate a RAM computation with an oblivious RAM where the data owner, Alice, has a stateful private memory of size $O(\sqrt{n})$, achieving an expected amortized time overhead of $O(\log^2 n)$ using $O(n \log n)$ memory at the data provider. In addition, Williams *et al.* [32] claim a method that uses an $O(\sqrt{n})$ -sized private cache and has $O(\log n \log \log n)$ amortized time overhead, but Pinkas and Reinman [26] have raised concerns with the assumptions and analysis of this result.

Goodrich and Mitzenmacher [16] provide a stateful RAM simulation scheme that achieves an overhead of $O(\log n)$ and is oblivious with very high probability. Their scheme assumes that Alice maintains state from one episode to the next in a private cache of size $O(n^\nu)$, for any given fixed constant $\nu > 0$. Boneh *et al.* [7] also propose a scheme that uses a state. They achieve an amortized overhead of $O(1)$ but using a state of size $O(\sqrt{n \log n})$. However, this state is essential to the efficiency of both simulation schemes. Thus, these methods are not applicable to the problem of providing privacy-preserving group access to an outsourced data repository.

Returning to stateless oblivious RAM simulation, we note that Ajtai [4] has a recent oblivious RAM simulation result that shows that a polylogarithmic factor overhead in time and space is possible without cryptographic assumptions about the existence of random hash functions, as is done in the previous oblivious RAM simulation cited above. Damgård *et al.* [10] improve this result further, showing that a time overhead of $O(\log^3 n)$ is possible for oblivious RAM simulation without using random functions.

In addition to the above-mentioned upper-bound results, Beame and Machmouchi [5] show that if the additional space utilized in the simulation (besides the space for the data itself) is sufficiently sublinear, then the overhead for oblivious RAM simulation has a superlogarithmic lower bound. Such bounds don’t

²The scheme of Pinkas and Reinman allows the adversary, Bob, to distinguish with high probability an access sequence that reads the same memory location over and over from one that accesses each memory cell exactly once. This flaw is expected to be repaired in the journal version of their paper.

Table 1: Comparison of Oblivious RAM simulations.

	User Memory	User State Size	Server Storage	Amortized Access Overhead
Goldreich and Ostrovsky [13]	$O(1)$	-	$O(n \log n)$	$O(\log^3 n)$
Williams and Sion [31]	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(n \log n)$	$O(\log^2 n)$
Goodrich and Mitzenmacher [16] (1)	$O(1)$	-	$O(n)$	$O(\log^2 n)$
Goodrich and Mitzenmacher [16] (2)	$O(n^\nu)$	$O(n^\nu)$	$O(n)$	$O(\log n)$
Boneh <i>et al.</i> [7]	$O(\sqrt{n \log n})$	$O(\sqrt{n \log n})$	$O(n)$	$O(1)$
Our result	$O(n^\nu)$	-	$O(n)$	$O(\log n)$
Our result (w/o random oracle)	$O(n^\nu)$	-	$O(n)$	$O(\log^2 n)$

apply, of course, to a simulation that uses $O(n)$ additional memory, as is common in the efficient schemes mentioned above.

We provide a summary of the Oblivious RAM simulation schemes and compare with ours in Table 1.3. Note that the schemes that maintain a state cannot be used to hide a pattern of access by a group of users which is one of the challenges we address in this paper.

1.4 Our Results

We give an efficient method for simulating any RAM algorithm, \mathcal{A} , in a stateless fashion with a time overhead of $O(\log n)$ and space overhead of $O(1)$, using an access sequence that is data-oblivious with very high probability, where n is the size of the RAM memory. Our methods assume that Alice has a private cache of size $O(n^\nu)$, for any given fixed constant $\nu > 0$, but she uses this cache only as a private “scratch space” to support computations she performs during each episode. Alice is not allowed to maintain state in her private memory from one episode to the next. Thus, this simulation scheme is applicable to the problem of simulating access to a shared data repository by a group of cooperating users that all share a secret key. Moreover, the assumption about the size of Alice’s scratch space is motivated by the fact that even handheld devices have a reasonable amount of local memory. For example, if we were to set $\nu = 1/4$, then our simulation would allow a collection of devices having memories with sizes on the order of one megabyte to support privacy-preserving access to an outsourced data repository whose size is on the order of one yottabyte.

Like the previous oblivious RAM simulation schemes mentioned above, our scheme uses a hierarchy of hash tables, together with a small set of pseudorandom hash functions, to obfuscate the access pattern of the algorithm \mathcal{A} (which need not be

specified in advance). The main idea of our scheme is to maintain these hash tables as cuckoo hash tables that all share a single stash of size $O(\log n)$. While conceptually simple, this approach requires a new, non-trivial analysis for a set of cuckoo tables sharing a common stash. In addition, an important technical detail that simplifies our construction is that we make no use of so-called “dummy” elements, whereas the previous schemes used such elements.

In practice, the set of pseudorandom hash functions could be implemented using, e.g., keyed SHA-256 functions [11]. Nevertheless, we also show that our construction can be used to simulate a RAM computation with an overhead of $O(\log^2 n)$ without the use of pseudorandom functions, which may be of some theoretical interest.

Finally, we provide experimental results for a simulation of our scheme, which show the practical effectiveness of the approach of using a shared stash. In particular, our experimental prototype simulates the dynamic evolution of the hierarchy of hash tables and our experimental analysis shows the threshold values at which the shared stash becomes effective.

2 Theory Background

For our results, we rely on general methods for data-oblivious simulation of a non-oblivious algorithm on a RAM. As mentioned above, the seminal theoretical framework for such simulations was presented by Goldreich and Ostrovsky [13], who store keys in a hierarchy of hash tables of increasing size, each being twice the size of the previous one. For n items there are $O(\log n)$ levels, each level being a standard hash table with 2^i buckets for some i , and each bucket containing up to $O(\log n)$ keys in order to cope with collisions within the hash table. In this construction the total size of all the tables is $O(n \log n)$. To perform a lookup, the first level is scanned sequentially,

and in each of the other levels, a bucket chosen by the hash function for that level acting on the key (or, if the item is found at an earlier level, a random dummy key) is scanned. The item is subsequently re-encrypted and re-inserted into the first level. It is important to note that at all levels a bucket is scanned even if the key is found early, to maintain obliviousness. As levels fill, keys must be shifted down to subsequent levels. The details of the original scheme are rather complex; for further details see the original paper [13].

Recently, a more efficient simulation approach for this problem was outlined by Goodrich and Mitzenmacher [16]. The primary difference in this new line of work is the use of *cuckoo hash tables* in place the standard hash tables used originally in [13]. We therefore now present some background on *cuckoo hashing*.

As introduced by Pagh and Rodler [24], in standard cuckoo hashing we utilize two tables, each with m cells, with each cell capable of holding a single key. We make use of two hash functions h_1 and h_2 . We assume that the hash functions can be modeled as completely random hash functions. The tables store up to n items, where $m = n(1 + \epsilon)$ for some constant $\epsilon > 0$, yielding a load of (just) less than $1/2$; keys can be inserted or deleted over time as long as this restriction is maintained.

A key x (which we may also refer to as an “item” or “element”) that is stored in the hash tables must be located at either $h_1(x)$ or $h_2(x)$. As there are only two possible locations for a key, lookups take constant time. To insert a new key x , we place x in the cell $h_1(x)$. If the cell had been empty, the operation is complete. Otherwise, key y previously in the cell is moved to $h_2(y)$. This may in turn require another key to be moved, and so on, until a key is placed in an empty cell. We say that a failure occurs if, for an appropriate constant c , after $c \log n$ steps this process has not successfully terminated. Suppose we insert an n th key into the system. It is known that:

- The expected time to insert a new key is bounded above by a constant (that depends on ϵ).
- The probability that a new key causes a failure is $\Theta(1/n^2)$ (where the notation hides a dependence on ϵ).

See Figures 1 and 2 for examples.

Before considering ways to reduce the probability of failures to something more suitable, we briefly mention that there are several natural variations of cuckoo hashing, many of which are described in a survey article by Mitzenmacher [22]. Variations include




Figure 1: The top of the figure represents a cuckoo hash table. Keys are placed in one subtable; the arrow for each key points to the alternate location for the key in the other subtable. Key G is inserted, leading to the movement of several other keys for G to be placed, as shown in the bottom of the figure.




Figure 2: Key G is to be inserted, but it cannot be placed successfully. (Seven keys have only six locations.) This leads to a failure, or if there is a stash, then G can be placed in a stash.

using more than two choices, using cells that hold more than one key, and so on. For our purposes, it suffices to understand standard cuckoo hashing, along with idea of a *stash* [17].

A stash represents additional memory where keys that would cause a failure can be placed in order to avoid the failure; with a stash, a failure occurs only if the stash itself overflows. As shown in [17], the failure probability when inserting the n th key into a cuckoo hash table can be reduced to $O(1/n^{k+2})$ for any constant k by using a stash that can hold k keys. Using this allows us to use cuckoo hash tables for any polynomially bounded number of inserts and deletions using only a constant-sized stash. To search for an item, we must search both the two table locations and the k stash locations. In the context of

oblivious simulation, we can search the stash simply by reading each stash location.

As we have stated, however, in order to perform our oblivious simulation, we will make use of a hierarchy of cuckoo hash tables to hold n items. The smallest of these hash tables may be much smaller than n , which can lead to problems in our setting³. For example, if the smallest hash table is of size x , then even using a stash of size k leads to a failure probability of $O(1/x^{k+2})$. If x is for example polylogarithmic in n , then for any constant k , the failure probability is $\Omega(1/n)$, and therefore over the insertion of n items, we would expect failures to occur. In order to deal with this problem, Goodrich and Mitzenmacher [16] extend the analysis of [17] to stashes of logarithmic size, showing that even for suitably large table sizes x that are only polylogarithmic in n , and stashes of size $k = O(\log n)$, the failure probability is $O(x^{-\alpha k})$ for a suitable constant k . This suffices to yield super-polynomially small failure rates.

In fact, we need to extend this result even further here. In [16], Goodrich and Mitzenmacher use a logarithmic-sized stash at each level. We explain here that it suffices to use a *single* logarithmic-sized stash for all levels. That is, while there's a non-trivial probability of *at least one* layer in our hierarchy requiring a stash of size $\Omega(1)$, over the logarithmic number of layers only a stash of logarithmic size is actually necessary. We will use this in our construction in Section 3. We now briefly explain why, if we consider the sum of the number of items placed in the stash at all possible levels in our construction, this will be at most $O(\log n)$ with high probability.

The key is the following argument. As shown in [16], at a level of size x cells (where x is $\Omega(\log^7 n)$), the probability that the stash for that level exceeds a total size s is $x^{-\Omega(s)}$. Further, as long as the hashes for each level in our construction are independent, we can treat the required stash size at each level is independent, since the number of items placed in the stash at a level is then a random variable dependent only on the number of items appearing in that level.

Now consider any point of our construction and let S_i be the number of items at the i th level that need to be put in the stash. It is apparent that S_i has mean less than 1 and tails that can be dominated by a geometrically decreasing random variable. This is sufficient to apply standard Chernoff bounds. Formally, let X_1, X_2, \dots, X_ℓ be independent random variables with mean 1 geometrically decreasing tails, so that $X_i = j$ with probability $1/2^j$ for $j \geq 1$. Then the calculations of [16] imply that the X_i stochasti-

cally dominate the S_i , and we can now apply standard Chernoff bounds for these random variables. Specifically, noting that X_i can be interpreted as the number of fair coin flips until the first heads, we can think of the sum of the X_i as being the number of coin flips until the ℓ th head, and this dominates the number of items that need to be placed in the stash at any point. When $\ell = O(\log n)$, as is the case here as there are only $O(\log n)$ levels of hash tables in our construction, then for any constant γ_1 there exists a corresponding constant γ_2 such that the ℓ th head occurs by the $(\gamma_2 \log n)$ 'th flip with probability at least $1 - 1/n^{\gamma_1}$. (See, for example, [23, Chapter 4].) Hence we can handle any polynomial number of steps with high probability, using a stash of size only $O(\log n)$ that holds items from all levels of our construction.

3 Simulating a RAM Algorithm Obliviously

In this section, we describe and analyze two schemes for stateless oblivious RAM simulation.

3.1 Simulation Using Pseudorandom Functions

We begin with a construction that uses pseudorandom functions and is secure against a polynomially bounded adversary.

Given a RAM algorithm, \mathcal{A} , the main goal of our oblivious simulation of \mathcal{A} is to hide the pattern of memory accesses that are made by \mathcal{A} . As mentioned in Section 2, we follow the general framework introduced by Goldreich and Ostrovsky [13], which uses a hierarchy of hash tables.

Let n be the number of memory cells of the RAM. We view each such cell as an item consisting of a pair (x, v) , where $x \in \{0, \dots, n-1\}$ is the index and v is the corresponding value. Our data structure stored at the server has three components, illustrated in Figure 3. The first component is a cache of size $O(\log n)$, denoted by Q . The second component is a hierarchy of cuckoo hash tables, $T = (T_1, \dots, T_L)$, where the size of T_1 is twice the size of Q , each table T_{i+1} is twice the size of table T_i , and T_L is the first table in the sequence of size greater than or equal to n . Thus, L is $O(\log n)$. The third component is a stash, S , shared between all the above cuckoo tables.

RAM items are stored in the data structure in encrypted form. We use a semantically secure probabilistic encryption scheme, which results in a different ciphertext for the same item each time it is re-encrypted. Also, the server is unable to determine

³This is also at the heart of the flaw in the CRYPTO2010 version of the Pinkas and Reinman paper [26].




Figure 3: Illustration of the data structure stored at the server for oblivious RAM simulation using pseudo-random functions. In the access phase of the simulation, all the items in the cache, Q and the stash, S , plus two items for each cuckoo table T_i are read by the server. The locations accessed by the server are visualized as gray-filled rectangles.

whether two ciphertexts correspond to the same item. The stash S is handled in a similar manner whenever we search in it for an item.

We use a family of pseudorandom functions parameterized by a secret value, k_i , for each table, T_i , such that no value k_i is revealed to the server. In particular, k_i is stored in encrypted form for each table T_i , so that each user can read k_i , decrypt it, and then use it to provide the two hash functions, h_1 and h_2 , employed by the cuckoo table, T_i , to determine the location of items. In particular, a memory item (x, v) is mapped to locations $h_1(x)$ and $h_2(x)$ in T_i by the cuckoo scheme (and stored in one of these two locations or in the common stash, S).

The data structure is initialized by storing all the n RAM items into cuckoo table T_L . Each memory access defined by algorithm \mathcal{A} corresponds to an *episode* in our simulation. An episode consists of two phases, an *access phase* and a *rebuild phase*.

Suppose algorithm \mathcal{A} calls for an access to memory item (x, v) . The access phase consists of a search for x in the cache, Q , then in the stash, S , and continues with a two-cell cuckoo lookup in each of T_1 to T_L until we find the first item with index x . Once we have found this item, we have achieved the goal of our search, of course. Nevertheless, for the sake of obliviousness, we simulate continuing the search throughout the entire data structure. Namely, we always traverse completely Q and S , and we perform two-cell cuckoo accesses in tables T_1 through T_L . However, after the item is found, we simply access two distinct, independent uniformly chosen random locations in each remaining cuckoo table.

Once we have completed the access phase, which takes $O(\log n)$ time, we then switch to the rebuild phase. We begin by adding or replacing a copy of the found item into cache Q , possibly changing its value in the case of a write operation. To assure obliviousness, we exhaustively scan Q in a sequential manner and re-encrypt and rewrite all its items. Thus, the server cannot distinguish which item was accessed and whether it was modified.

We note briefly that if the item is in the stash, we can obviously remove it from the stash when placing it into Q , to help make sure the stash does not overflow. One natural approach is to have stash cells have an associated “clean” or “dirty” bit, which is encrypted along with the rest of the item. A clean cell can store an item; a dirty cell is currently being utilized. When an item is found and replaced into Q , we can set the cell to clean in the stash.

After adding enough items, cache Q will eventually overflow. We remedy the overflow by moving all the elements of Q to cuckoo table T_1 , including those associated with empty locations. However, in order to maintain obliviousness, we do not wait for an overflow to occur and instead perform the move after a number of accesses equal to the size of Q . The moving down of elements cascades down through the hierarchy of cuckoo tables at a fixed schedule by periodically moving the elements of T_{i-1} into T_i at the earliest time T_{i-1} could have become full. Also, suppose that we are going to move elements into table T_i for the second time, then we instead move the elements into table T_{i+1} . Moreover, we continue applying this rule for $i = 1, 2, \dots$, until we are copying the elements

into a table for the first time or we reach T_L . Thus, the process of copying elements into a cuckoo table occurs at deterministic instances, depending only on the place we currently are at in the access sequence specified by algorithm \mathcal{A} .

In order to move m elements from a table T_i into a cuckoo hash table T_{i+1} obviously, we use an algorithm of [16] to obviously sort the items using $O(m)$ accesses to the outsourced memory, assuming we have a private workspace of size $O(n^\nu)$, for some constant $\nu > 0$, and $m \geq \log n$, which is always true in our case. This allows us to remove duplicate items and use another algorithm of [16] to obviously construct a cuckoo table of size m and an associated stash, S' , of size $O(\log n)$ in $O(m)$ time, with very high probability, while utilizing the private workspace of size $O(n^\nu)$. Given this construction, we then read S and S' into our private workspace, remove any duplicates and merge them into a single stash S (which will succeed with very high probability, based on the analysis we have given above), and write S back out in a straightforward oblivious fashion. Note that in order to assure obliviousness in subsequent lookups, table T_{i+1} is rebuilt using two new pseudorandom hash functions selected by the client by replacing parameter k_{i+1} with a new secret value.

Any access performed in our simulation will eventually lead to $O(\log n)$ table rebuilds, with each element in a rebuild being charged with a constant amount of work; hence, the total amortized overhead of all rebuild phases is $O(\log n)$. Therefore, the total amortized time overhead of the entire simulation is $O(\log n)$. Moreover, it is easy to see that the space used by the data structure stored at the server is $O(n)$.

Let us therefore consider the obliviousness of this scheme. As we have already observed, the rebuild phase is clearly oblivious, so any potential dependencies on input values would have to come in the access phase. Recall that in the access phase, we search in S , Q , and do a two-table cuckoo access in T_1, \dots, T_L . Moreover, because we move the found item into Q after each access, and we switch to performing random table lookups once we have found the item, we are guaranteed never to repeat a two-cell cuckoo lookup in any table, T_i , for the same item x . In addition, each such lookup is an independent uniformly random access to a table (either from our assumption about h_1 and h_2 being distinct random hash functions for each table or because we already found the item and are making random accesses explicitly). We perform $O(|T_i|)$ such lookups before we empty T_i ; hence, the obliviousness of our access sequence depends on the inability of the adversary, Bob, of telling if we

are doing a search for an actual item or performing a random access for the sake of obliviousness. That is, with high probability, Bob should not be able to determine whether the item was in S , Q , or some T_i at the point we found it. Note that this ability depends solely on whether or not the $O(|T_i|)$ accesses we made to T_i , together with searches in the shared stash S , would correspond to valid cuckoo lookups in T_i for some set of items. Of course, this is the same as the event that inserting all these elements into T_i would form a valid cuckoo table, with shared stash S , which we have already observed (in Section 2) is an event that occurs with very high probability. Thus, our scheme is oblivious with very high probability.

3.2 Simulation Without Pseudorandom Functions

We can adapt our simulation to avoid the use of random functions by employing an elegant trick due to Damgård *et al.* [10], albeit now further simplified to avoid the use of dummy nodes, which would add an extra level of complication that our scheme doesn't require.

The main idea is to place a complete binary tree, B , on top of all the memory cells used in the algorithm \mathcal{A} , and access each memory cell x by performing a binary search from the root of B to the leaf node corresponding to x . That is, we associate each memory cell item used by \mathcal{A} with a leaf of B , define B to have height $\lceil \log n \rceil$, and include information at each internal node v of B so that a search for x can determine in $O(1)$ time whether to proceed with the left child or right child of v . In our case, we store each of the nodes of B in our hierarchy of tables, similar to what is described above, with the shared stash, S , the cache, Q , and the set of cuckoo tables, T_1 to T_L . (See Figure 4.)

The main difference of this scheme with that given above is that in this case we no longer use random hash functions, h_1 and h_2 , to determine the locations of each element x in a cuckoo table T_i . Instead, we simply choose two distinct, independent uniformly random locations, i_1 and i_2 , in the respective two sides of T_i and associate these with x as a tuple (x, i_1, i_2) , which now represents the element x in our table.

Initially, all the nodes of B are stored in this way in T_L , and for each such internal node v , we include in v 's record pointers to the two random indices (and table index) for v 's left child and pointer to the two random indices (and table index) for v 's right child. Such pointers can be built obliviously by $O(1)$ calls to oblivious sorting once we have placed all the nodes




Figure 4: Illustration of the data structure stored at the server for oblivious RAM simulation without using pseudorandom functions. The binary tree is shown conceptually on the right and in terms of the storage locations of its nodes on the left. The storage locations for the nodes of the binary tree, B , are visualized as gray-filled rectangles. In the access phase of the simulation for a non-root node, all the items in the cache, Q and the stash, S , plus two items for each cuckoo table T_i are read by the server.

into T_L . Moreover, we will maintain such pointers throughout our simulation. In addition, we store the root r of B separately, as it is accessed in every step of our simulation.

Let us consider, therefore, how an access now occurs. The critical property, which we maintain inductively, is that, for each node v in B , which, say, is stored in T_j as its earliest (highest) location in our hierarchy, all the ancestors of B are stored in the tables T_1, \dots, T_j , or in r , S , or Q .

Our access for a memory cell x now occurs as a root-to-leaf search in B . We begin by searching in r to identify the two random indices and the table index for each of r 's children. Based on the value of x , we need to search next for either the left or right child of r , so let i_1 and i_2 be the two random indices for this node, w , and let j be the index of the highest table T_j storing w (with $j = 0$ if $w \in Q$ and $j = -1$ if $w \in S$). We next search in S and Q for w , and then proceed in T_1 through T_L . Of course, we already know the table where we will find w . So, for each table T_k with $k \neq j$, we simply access two random locations in T_k for the sake of obliviousness. For T_j itself, we look in locations $T_j[i_1]$ and $T_j[i_2]$ to find the cell containing the record for w . If w is not a leaf node, we repeat the above lookup search for the appropriate child of w that will lead us to the node storing x .

Once we have done our lookup for x , and have accessed a root-to-leaf set of nodes,

$$W = \{w_1, w_2, \dots, w_{\log n}\},$$

in the process, we perform a rebuild phase for W , as in the above construction based on random hash functions, except that we use random locations for all the nodes we move rather than use random functions. Note that by our induction hypothesis, if we move a set of nodes into a table T_i , then all the pointers for these nodes are either in T_i itself (hence, can be identified after $O(1)$ calls to oblivious sorting, which takes $O(n)$ memory accesses by the algorithm of [16]) or at lower levels in the hierarchy (hence, these pointers don't change by our move into T_i). Moreover, all the nodes of W move as a group. Thus, any root-to-leaf path in B must be stored in the tables T_1 to T_L , plus the queue Q and stash S , in a way that satisfies our induction hypothesis.

The lookup for an element x now requires searching for $O(\log n)$ nodes of B in our hierarchy, which costs an amortized overhead of $O(\log n)$ time each. Thus, each lookup costs us an amortized overhead of $O(\log^2 n)$ time. The obliviousness of this simulation follows from an argument similar to that given above for the obliviousness for our method that uses random hash functions. Therefore, we can perform a stateless oblivious RAM simulation without using random hash functions with an amortized time overhead of $O(\log^2 n)$, assuming a private workspace of size $O(n^\nu)$ for some constant $\nu > 0$.

4 Performance

We have implemented a preliminary prototype of our method for oblivious RAM simulation based on pseudorandom functions (Section 3.1) with the goal of estimating the size of the stash S needed to avoid failures during the rebuild phase. A failure can happen when we move elements from table T_i to T_{i+1} and the stash overflows, in which case we need to rebuild table T_{i+1} . In this section we present experimental results and show that for a small constant s , a stash of size $s \log n$ is enough to avoid failures.

Our prototype simulates the dynamic evolution of the hierarchy of hash tables during the access and rebuild phases, omitting the steps that maintain obliviousness (e.g. copying the stash to the client’s side). We maintain a stash S , a cache Q of size $\log n$ and a hierarchy of $O(\log n)$ cuckoo hash tables T_1 to T_L , where L is the smallest i such that $2^i \log n \geq n$. T_i consists of two hash tables of size $(1 + \epsilon)2^i \log n$ with hash functions h_1^i and h_2^i . Every memory access is followed by the insertion of the corresponding element into Q . If the item was retrieved from stash S , it is first copied to Q and then is removed from S . We move all the elements of Q to T_1 when the number of performed accesses is a multiple of $\log n$. Similarly, we move all the elements of T_i to T_{i+1} when the number of performed accesses is a multiple of $2^i \log n$. New hash functions are picked for both tables during this phase. During the insertion, an item is placed into the stash if after $c \log n$ moves it has not found an empty cell in the table; we experiment using $c = 2$. We insert an item into the stash only if it is not already present there.

Our prototype is implemented in Java. To generate hash functions we use a variation of a method recommended in [11], where $h_1^i(x) = \text{SHA256}(x \parallel \text{seed}_i^1) \bmod n$, and similarly for h_2^i . The seeds are 64-bit long and were obtained using a SHA256 hash chain starting from an initial seed.

We emphasize that in any implementation of our method there are various tradeoffs. For example, increasing the space (that is, using larger values of ϵ) reduces the average time for an insertion and the failure probability, as it reduces the frequency with which items have to be put in the stash. Increasing the stash size reduces the failure probability at the expense of additional time to examine the stash at each step. Increasing the number of moves allowed before placing an item in the stash increases the time but lowers the failure probability. Our purpose here is not to explore this broad range of tradeoffs, but to demonstrate the feasibility of this approach; exploring finer tradeoffs is left as future work.

We ran our simulation for up to 1024K ($K = 1000$) RAM items and a varying number of requests. Our experiments use a value of ϵ of 0.1 and 0.2. For each experiment we recorded the lowest size of S that is needed to avoid a failure. In Figures 5 and 6 we show the fraction of trials out of 1000 that result in the stash overflow. Comparing the two figures, we see that overflows happen substantially more frequently with $\epsilon = 0.1$ than with $\epsilon = 0.2$, as one would expect since smaller tables lead to more collisions. Indeed, for $\epsilon = 0.2$ we found a stash size of less than $\log n$ was enough to avoid overflows completely in our limited experiments. Also, a higher number of requests requires a slightly bigger stash since rebuilding happens more often, leading to a larger maximum stash requirement. While much more extensive experimentation would be needed to determine suitable stash sizes that would avoid stash overflow for numbers of trials many orders of magnitude larger, recall that the probability of an additional item needing to be placed in a stash in a standard cuckoo table falls very quickly. We thus expect only slightly larger stash sizes for such improvements in robustness.




Figure 5: Failure rate in 1000 trials for n items and r requests with $c = 2$, $\epsilon = 0.1$.

5 Conclusion

We have given schemes for achieving privacy-preserving access, with very high probability, to an outsourced data repository for a group of trusted users. Our scheme assumes each user has a modest amount of private memory, of size $O(n^\nu)$, for any given fixed constant $\nu > 0$, which is used as a workspace for private computation and carries no




Figure 6: Failure rate in 1000 trials for n items and r requests with $c = 2$, $\epsilon = 0.2$.

state from one interaction with the data repository to the next. Assuming the existence of pseudorandom hash functions, say implemented as keyed SHA-256 functions in practice, our protocol has an $O(\log n)$ amortized time overhead and an $O(1)$ space overhead. Moreover, our experiments show that this protocol would be effective in practice. If pseudorandom hash functions are not to be used, then we show that our protocol can be adapted to have an overhead of $O(\log^2 n)$.

There are several directions for future work, including the following:

- Our protocols assume that the manager of the data repository, Bob, is honest-but-curious.
 - Can our schemes be efficiently adapted to the case where Bob is only semi-trusted?
 - Can we efficiently handle a situation where Bob acts maliciously against some users?
 - Can we prevent Bob from performing a replay attack on some users using old versions of his memory?
- Our protocols also assume that the members of the group of cooperating users are all trusted.
 - What if some of the members of the group are malicious?
 - What if some of them collude with Bob to try to reveal the access patterns of other users?

Acknowledgments

This research was supported in part by the National Science Foundation under grants 0724806, 0713046, 0847968, 0953071, and 1012060.

References

- [1] A. V. Aho, J. E. Hopcroft, and J. D. Ullman, *Data Structures and Algorithms*. Reading, MA: Addison-Wesley, 1983.
- [2] M. Ajtai, J. Komlós, and E. Szemerédi, “An $O(n \log n)$ sorting network,” in *Proc. ACM Symposium on Theory of Computing (STOC)*, 1983, pp. 1–9.
- [3] —, “Sorting in $c \log n$ parallel steps,” *Combinatorica*, vol. 3, pp. 1–19, 1983.
- [4] M. Ajtai, “Oblivious RAMs without cryptographic assumptions,” in *Proc. ACM Symposium on Theory of Computing (STOC)*. ACM, 2010, pp. 181–190.
- [5] P. Beame and W. Machmouchi, “Making RAMs oblivious requires superlogarithmic overhead,” *Electronic Colloquium on Computational Complexity*, Report TR10-104, 2010, <http://eccc.hpi-web.de/report/2010/104/>.
- [6] M. Bellare and P. Rogaway, “Random oracles are practical: a paradigm for designing efficient protocols,” in *Proc. ACM Conference on Computer and Communications Security (CCS)*. New York, NY, USA: ACM, 1993, pp. 62–73.
- [7] D. Boneh, D. Mazieres, and R. A. Popa, “Remote oblivious storage: Making oblivious RAM practical,” Technical Report, 2011, <http://dSPACE.mit.edu/handle/1721.1/62006>.
- [8] S. Chen, R. Wang, X. Wang, and K. Zhang, “Side-channel leaks in web applications: a reality today, a challenge tomorrow,” in *Proc. IEEE Symposium on Security and Privacy*, 2010, pp. 191–206.
- [9] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. Cambridge, MA: MIT Press, 2001.
- [10] I. Damgård, S. Meldgaard, and J. B. Nielsen, “Perfectly secure oblivious RAM without random oracles,” *Cryptology ePrint Archive*, Report 2010/108, 2010, <http://eprint.iacr.org/>.

- [11] Y. Dodis and P. Puniya, “Getting the best out of existing hash functions; or what if we are stuck with SHA?” in *Applied Cryptography and Network Security (ACNS)*, 2008, pp. 156–173.
- [12] O. Goldreich, S. Micali, and A. Wigderson, “How to play ANY mental game,” in *Proc. ACM Symposium on Theory of Computing (STOC)*. New York, NY, USA: ACM, 1987, pp. 218–229.
- [13] O. Goldreich and R. Ostrovsky, “Software protection and simulation on oblivious RAMs,” *J. ACM*, vol. 43, no. 3, pp. 431–473, 1996.
- [14] M. T. Goodrich and R. Tamassia, *Algorithm Design: Foundations, Analysis, and Internet Examples*. New York, NY: John Wiley & Sons, 2002.
- [15] M. T. Goodrich, “Randomized Shellsort: A simple oblivious sorting algorithm,” in *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2010, pp. 1–16.
- [16] M. T. Goodrich and M. Mitzenmacher, “Privacy-preserving access of outsourced data via oblivious RAM simulation,” in *Proceedings of ICALP*, 2011, to appear.
- [17] A. Kirsch, M. Mitzenmacher, and U. Wieder, “More robust hashing: cuckoo hashing with a stash,” *SIAM J. Comput.*, vol. 39, pp. 1543–1561, 2009.
- [18] J. Kleinberg and E. Tardos, *Algorithm Design*. Boston, MA, USA: Addison-Wesley, Inc., 2005.
- [19] D. E. Knuth, *Sorting and Searching*, ser. The Art of Computer Programming. Reading, MA: Addison-Wesley, 1973, vol. 3.
- [20] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*. San Mateo, CA: Morgan-Kaufmann, 1992.
- [21] T. Leighton and C. G. Plaxton, “Hypercube sorting networks,” *SIAM J. Comput.*, vol. 27, no. 1, pp. 1–47, 1998.
- [22] M. Mitzenmacher, “Some open questions related to cuckoo hashing,” in *Proc. European Symposium on Algorithms (ESA)*, 2009, pp. 1–10.
- [23] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. New York, NY, USA: Cambridge University Press, 2005.
- [24] R. Pagh and F. Rodler, “Cuckoo hashing,” *Journal of Algorithms*, vol. 52, pp. 122–144, 2004.
- [25] M. Paterson, “Improved sorting networks with $O(\log N)$ depth,” *Algorithmica*, vol. 5, no. 1, pp. 75–92, 1990.
- [26] B. Pinkas and T. Reinman, “Oblivious RAM revisited,” in *Advances in Cryptology (CRYPTO)*, ser. Lecture Notes in Computer Science, T. Rabin, Ed. Springer, 2010, vol. 6223, pp. 502–519.
- [27] N. Pippenger and M. J. Fischer, “Relations among complexity measures,” *J. ACM*, vol. 26, no. 2, pp. 361–381, 1979.
- [28] V. R. Pratt, “Shellsort and sorting networks,” Ph.D. dissertation, Stanford University, Stanford, CA, USA, 1972.
- [29] J. Seiferas, “Sorting networks of logarithmic depth, further simplified,” *Algorithmica*, vol. 53, no. 3, pp. 374–384, 2009.
- [30] G. Wang, T. Luo, M. T. Goodrich, W. Du, and Z. Zhu, “Bureaucratic protocols for secure two-party sorting, selection, and permuting,” in *Proc. ACM Symposium on Information, Computer and Communications Security (ASIACCS)*. New York, NY, USA: ACM, 2010, pp. 226–237.
- [31] P. Williams and R. Sion, “Usable PIR,” in *NDSS*. The Internet Society, 2008.
- [32] P. Williams, R. Sion, and B. Carbunar, “Building castles out of mud: practical access pattern privacy and correctness on untrusted storage,” in *Proc. ACM Conference on Computer and Communications Security (CCS)*. New York, NY, USA: ACM, 2008, pp. 139–148.